



> home > about > feedback > login

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: [optimal <and> range <and> presentation <and> hierarchy <and> sub-range<AND>((optimal <and> range <and> presentation <and> hierarchy))]

Found 12 of 151,219 searched.

Search within Results



> Advanced Search

> Search Help/Tips

Sort by: Title Publication Publication Date Score Binder

Results 1 - 12 of 12 short listing

- 1** Invited talk: Managing dynamic concurrent tasks in embedded real-time 80%
 multimedia systems

Peng Yang , Paul Marchal , Chun Wong , Stefaan Himpe , Francky Catthoor , Patrick David , Johan Vounckx , Rudy Lauwereins

Proceedings of the 15th international symposium on System Synthesis October 2002

This paper addresses the problem of mapping an application, which is highly dynamic in the future, onto a heterogeneous multiprocessor platform in an energy efficient way. A two-phase scheduling method is used for that purpose. By exploring the Pareto curves and scenarios generated at design time, the run-time scheduler can easily find a good scheduling at a very low overhead, satisfying the system constraints and minimizing the energy consumption. A real-life example from a 3D quality of service ...

- 2** Range queries in OLAP data cubes 80%

Ching-Tien Ho , Rakesh Agrawal , Nimrod Megiddo , Ramakrishnan Srikant
ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data June 1997

Volume 26 Issue 2

A range query applies an aggregation operation over all selected cells of an OLAP data cube where the selection is specified by providing ranges of values for numeric dimensions. We present fast algorithms for range queries for two types of aggregation operations: SUM and MAX. These two operations cover techniques required for most popular aggregation operations, such as those supported by SQL. For range-sum queries, the essential idea is to precompute some auxiliary information ...

- 3** File servers for network-based distributed systems 77%

Liba Svobodova
ACM Computing Surveys (CSUR) December 1984
 Volume 16 Issue 4

- 4 Energy-aware systems: Binary translation to improve energy efficiency through post-pass register re-allocation** 77%

 Kun Zhang , Tao Zhang , Santosh Pande

Proceedings of the fourth ACM international conference on Embedded software
September 2004

Energy efficiency is rapidly becoming a first class optimization parameter for modern systems. Caches are critical to the overall performance and thus, modern processors (both high and low-end) tend to deploy a cache with large size and high degree of associativity. Due a large size cache power takes up a significant percentage of total system power. One important way to reduce cache power consumption is to reduce the dynamic activities in the cache by reducing the dynamic load-store counts. In ...

- 5 Database selection for processing k nearest neighbors queries in distributed environments** 77%

 Clement Yu , Prasoon Sharma , Weiyi Meng , Yan Qin

Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries January 2001

We consider the processing of digital library queries, consisting of a text component and a structured component in distributed environments. The text component can be processed using techniques given in previous papers such as [7, 8, 11]. In this paper, we concentrate on the processing of the structured component of a distributed query. Histograms are constructed and algorithms are given to provide estimates of the desirabilities of the databases with respect to the given query. Databases ...

- 6 Construction of optimal graphs for bit-vector compression** 77%

 A. Bookstein , S. T. Klein

Proceedings of the 13th annual international ACM SIGIR conference on Research and development in information retrieval December 1989

Bitmaps are data structures occurring often in information retrieval. They are useful; they are also large and expensive to store. For this reason, considerable effort has been devoted to finding techniques for compressing them. These techniques are most effective for sparse bitmaps. We propose a preprocessing stage, in which bitmaps are first clustered and the clusters used to transform their member bitmaps into sparser ones, that can be more effectively compressed. The clustering method e ...

- 7 Virtual memory and backing storage management in multiprocessor operating systems using object-oriented design techniques** 77%

 V. F. Russo , R. H. Campbell

ACM SIGPLAN Notices , Conference proceedings on Object-oriented programming systems, languages and applications September 1989

Volume 24 Issue 10

The Choices operating system architecture [3, 4, 15] uses class hierarchies and object-oriented programming to facilitate the construction of customized operating systems for shared memory and networked multiprocessors. The software is being used in the Tapestry Parallel Computing Laboratory at the University of Illinois to study the performance of algorithms, mechanisms, and policies for parallel systems. This paper describes the architectural design and class hierarchy of ...

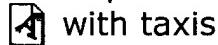
- 8 Research papers: data mining: An integrated approach for scaling up classification and prediction algorithms for data mining** 77%

 Patricia E. N. Lutu

Proceedings of the 2002 annual research conference of the South African institute of computer scientists and information technologists on Enablement through technology September 2002

Classification and prediction algorithms for machine learning typically require all training data to be resident in memory during decision tree construction. Typically, a flat file is created from database or data warehouse data and loaded into memory for processing. This severely limits the scalability of these algorithms to practical data mining tasks. Some attempts have been made by researchers to implement disk-based algorithms which can handle much larger training sets. Both approaches suff ...

9 Implementation of a compiler for a semantic data model: Experiences with taxis 77%



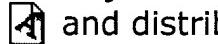
Brian Nixon , Lawrence Chung , John Mylopoulos , David Lauzon , Alex Borgida , M. Stanley

ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data December 1987

Volume 16 Issue 3

The features of a compiler for the Taxis design language are described and discussed. Taxis offers an entity-based framework for designing interactive information systems and supports generalisation, classification and aggregation as abstraction mechanisms. Its features include multiple inheritance of attributes, isA hierarchies of transactions, metaclasses, typed attributes, a procedural exception-handling mechanism and an iteration construct based on the abstraction mechanisms supported D ...

10 Object views: language support for intelligent object caching in parallel and distributed computations 77%



Ilya Lipkind , Igor Pechtchanski , Vijay Karamcheti

ACM SIGPLAN Notices , Proceedings of the 14th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications October 1999

Volume 34 Issue 10

Object-based parallel and distributed applications are becoming increasingly popular, driven by the programmability advantages of component technology and a flat shared-object space. However, the flat shared-object space introduces a performance challenge: applications that rely on the transparent coherent caching of objects achieve high performance only on tightly coupled parallel machines. In distributed environments, the overheads of object caching force application designers to choose o ...

11 On defining application-specific high-level array operations by means of shape-invariant programming facilities 77%



Sven-Bodo Scholz

ACM SIGPLAN APL Quote Quad , Proceedings of the APL98 conference on Array processing language July 1998

Volume 29 Issue 3

Most of the existing high-level array-processing languages support a fixed set of pre-defined array *operations* and a few higher-order functions for constructing new array operations from existing ones. In this paper, we discuss a more general approach made feasible by SAC (for Single Assignment C), a functional variant of C. SAC provides a meta-level language construct called WITH-loop which may be considered a sophisticated variant of the FORALL-loops ...

12 The MASC/BGMP architecture for inter-domain multicast routing 77%



Satish Kumar , Pavlin Radoslavov , David Thaler , Cengiz Alaettinoğlu , Deborah Estrin ,
Mark Handley
**ACM SIGCOMM Computer Communication Review , Proceedings of the ACM
SIGCOMM '98 conference on Applications, technologies, architectures, and
protocols for computer communication October 1998**

Volume 28 Issue 4

Multicast routing enables efficient data distribution to multiple recipients. However, existing work has concentrated on extending single-domain techniques to wide-area networks, rather than providing mechanisms to realize inter-domain multicast on a global scale in the Internet. We describe an architecture for inter-domain multicast routing that consists of two complementary protocols. The Multicast Address-Set Claim (MASC) protocol forms the basis for a hierarchical address allocation architect ...

Results 1 - 12 of 12 short listing

The ACM Portal is published by the Association for Computing Machinery. Copyright ?2005 ACM, Inc.



> home > about > feedback > login

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Nothing Found

Your search for the **Phrase optimal <and> range <and> presentation <and> hierarchy <and> truncate <and> divide <and> wildcard** did not return any results.
To search for *terms* separate them with **AND** or **OR**.

Click on the suggested options:

optimal AND <and> AND range AND <and> AND presentation AND <and> AND hierarchy AND <and> AND truncate AND <and> AND divide AND <and> AND wildcard

optimal OR OR range OR OR presentation OR OR hierarchy OR OR truncate OR OR divide OR OR wildcard

To search for names try using only the last or first name.

You may revise it and try your search again below or click advanced search for more options.

optimal <and> range <and>
 presentation <and> hierarchy
 <and> truncate <and> divide
 <and> wildcard



[\[Advanced Search\]](#) [\[Search Help/Tips\]](#)



Complete Search Help and Tips

The following characters have specialized meaning:

Special Characters	Description
, () [These characters end a text token.
= > < !	These characters end a text token because they signify the start of a field operator. (! is special: != ends a token.)
' @ \Q < { [!	These characters signify the start of a delimited token. These are terminated by the end character associated with the start character.



> home > about > feedback > login

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: [optimal <and> range <and> presentation <and> hierarchy <and> truncate <and> divide <and> magnitude]
Found 67 of 151,219 searched.

Search within Results



> Advanced Search

> Search Help/Tips

Sort by: Title Publication Publication Date Score Binder

Results 1 - 20 of 67 short listing

Prev Page 1 2 3 4
 Next Page

- 1** SMARTS: accelerating microarchitecture simulation via rigorous statistical sampling 80%

Roland E. Wunderlich , Thomas F. Wenisch , Babak Falsafi , James C. Hoe
ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture May 2003

Volume 31 Issue 2

Current software-based microarchitecture simulators are many orders of magnitude slower than the hardware they simulate. Hence, most microarchitecture design studies draw their conclusions from drastically truncated benchmark simulations that are often inaccurate and misleading. This paper presents the Sampling Microarchitecture Simulation (SMARTS) framework as an approach to enable fast and accurate performance measurements of full-length benchmarks. SMARTS accelerates simulation by selectively ...

- 2** Curriculum 68: Recommendations for academic programs in computer science: a report of the ACM curriculum committee on computer science 80%

William F. Atchison , Samuel D. Conte , John W. Hamblen , Thomas E. Hull , Thomas A. Keenan , William B. Kehl , Edward J. McCluskey , Silvio O. Navarro , Werner C. Rheinboldt , Earl J. Schweppe , William Viavant , David M. Young

Communications of the ACM March 1968

Volume 11 Issue 3

- 3** The Quadtree and Related Hierarchical Data Structures 80%

Hanan Samet
ACM Computing Surveys (CSUR) June 1984

Volume 16 Issue 2

4 An efficient I/O interface for optical disks 77%

Jeffrey S. Vitter

ACM Transactions on Database Systems (TODS) June 1985

Volume 10 Issue 2

We introduce the notion of an I/O interface for optical digital (write-once) disks, which is quite different from earlier research. The purpose of an I/O interface is to allow existing operating systems and application programs that use magnetic disks to use optical disks instead, with minimal change. We define what it means for an I/O interface to be disk-efficient. We demonstrate a practical disk-efficient I/O interface and show that its I/O performance in many cases is optimum, up to a ...

5 Practical byzantine fault tolerance and proactive recovery 77%

Miguel Castro , Barbara Liskov

ACM Transactions on Computer Systems (TOCS) November 2002

Volume 20 Issue 4

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ...

6 Gross motion planning—a survey 77%

Yong K. Hwang , Narendra Ahuja

ACM Computing Surveys (CSUR) September 1992

Volume 24 Issue 3

Motion planning is one of the most important areas of robotics research. The complexity of the motion-planning problem has hindered the development of practical algorithms. This paper surveys the work on gross-motion planning, including motion planners for point robots, rigid robots, and manipulators in stationary, time-varying, constrained, and movable-object environments. The general issues in motion planning are explained. Recent approaches and their performances are briefly described, a ...

7 Routing with a clue 77%

Yehuda Afek , Anat Bremler-Barr , Sariel Har-Peled

IEEE/ACM Transactions on Networking (TON) December 2001

Volume 9 Issue 6

We suggest a new simple forwarding technique to speed up IP destination address lookup. The technique is a natural extension of IP, requires 5 bits in the IP header (IPv4, 7 in IPv6), and performs IP lookup nearly as fast as IP/Tag switching but with a smaller memory requirement and a much simpler protocol. The basic idea is that each router adds a "clue" to each packet, telling its downstream router where it ended the IP lookup. Since the forwarding tables of neighboring routers are similar, a ...

8 Information systems outsourcing: a survey and analysis of the literature 77%

Jens Dibbern , Tim Goles , Rudy Hirschheim , Bandula Jayatilaka

ACM SIGMIS Database November 2004

Volume 35 Issue 4

In the last fifteen years, academic research on information systems (IS) outsourcing has evolved rapidly. Indeed the field of outsourcing research has grown so fast that there has been scant opportunity for the research community to take a collective

breath, and complete a global assessment of research activities to date. This paper seeks to address this need by exploring and synthesizing the academic literature on IS outsourcing. It offers a roadmap of the IS outsourcing literature, highlight ...

9 Object-based and image-based object representations

77%



Hanan Samet

ACM Computing Surveys (CSUR) June 2004

Volume 36 Issue 2

An overview is presented of object-based and image-based representations of objects by their interiors. The representations are distinguished by the manner in which they can be used to answer two fundamental queries in database applications: (1) Feature query: given an object, determine its constituent cells (i.e., their locations in space). (2) Location query: given a cell (i.e., a location in space), determine the identity of the object (or objects) of which it is a member as well as the re ...

10 Topical web crawlers: Evaluating adaptive algorithms

77%



Filippo Menczer , Gautam Pant , Padmini Srinivasan

ACM Transactions on Internet Technology (TOIT) November 2004

Volume 4 Issue 4

Topical crawlers are increasingly seen as a way to address the scalability limitations of universal search engines, by distributing the crawling process across users, queries, or even client computers. The context available to such crawlers can guide the navigation of links with the goal of efficiently locating highly relevant target pages. We developed a framework to fairly evaluate topical crawling algorithms under a number of performance metrics. Such a framework is employed here to evalua ...

11 Statistical Timing Analysis for Intra-Die Process Variations with Spatial Correlations

77%



Aseem Agarwal , David Blaauw , Vladimir Zolotov

Proceedings of the 2003 IEEE/ACM international conference on Computer-aided design November 2003

Process variations have become a critical issue in performance verification of high-performance designs. We present a new, statistical timing analysis method that accounts for inter- and intra-die process variations and their spatial correlations. Since statistical timing analysis has an exponential run time complexity, we propose a method whereby a statistical bound on the probability distribution function of the exact circuit delay is computed with linear run time. First, we develop a model for repre ...

12 XML query processing I: Dynamic sample selection for approximate query processing

77%



Brian Babcock , Surajit Chaudhuri , Gautam Das

Proceedings of the 2003 ACM SIGMOD international conference on Management of data June 2003

In decision support applications, the ability to provide fast approximate answers to aggregation queries is desirable. One commonly-used technique for approximate query answering is sampling. For many aggregation queries, appropriately constructed biased (non-uniform) samples can provide more accurate approximations than a uniform sample. The optimal type of bias, however, varies from query to query. In this paper, we describe an approximate query processing technique that dynamically constructs ...

- 13 Lexicon design using perfect hash functions** 77%
 Nick Cercone , Max Krause , John Boates
ACM SIGSOC Bulletin , Proceedings of the joint conference on Easier and more productive use of computer systems. (Part - II): Human interface and the user interface - Volume 1981 May 1981
Volume 13 Issue 2-3
The research reported in this paper derives from the recent algorithm of Cichelli (1980) for computing machine-independent, minimal perfect hash functions of the form: hash value ≡ hash key length + associated value of the key's first letter + associated value of the key's last letter hash value ≡ hash key length + associated value of the key's first letter + associated value of the key's last letter A minimal perfect hash function is one which provides sin ...
- 14 The early history of COBOL** 77%
 Jean E. Sammet
ACM SIGPLAN Notices , The first ACM SIGPLAN conference on History of programming languages January 1978
Volume 13 Issue 8
This paper discusses the early history of COBOL, starting with the May 1959 meeting in the Pentagon which established the Short Range Committee which defined the initial version of COBOL, and continuing through the creation of COBOL 61. The paper gives a detailed description of the committee activities leading to the publication of the first official version, namely COBOL 60. The major inputs to COBOL are discussed, and there is also a description of how and why some of the technical decisi ...
- 15 Scalable feature selection, classification and signature generation for organizing large text databases into hierarchical topic taxonomies** 77%
 Soumen Chakrabarti , Byron Dom , Rakesh Agrawal , Prabhakar Raghavan
The VLDB Journal — The International Journal on Very Large Data Bases August 1998
Volume 7 Issue 3
We explore how to organize large text databases hierarchically by topic to aid better searching, browsing and filtering. Many corpora, such as internet directories, digital libraries, and patent databases are manually organized into topic hierarchies, also called *taxonomies*. Similar to indices for relational data, taxonomies make search and access more efficient. However, the exponential growth in the volume of on-line textual information makes it nearly impossible to maintain such taxono ...
- 16 A radar simulation program for a 1024-processor hypercube** 77%
 J. L. Gustafson , R. E. Benner , M. P. Sears , T. D. Sullivan
Proceedings of the 1989 ACM/IEEE conference on Supercomputing August 1989
We have developed a fast parallel version of an existing synthetic aperture radar (SAR) simulation program, SRIM. On a 1024-processor NCUBE hypercube it runs an order of magnitude faster than on a CRAY X-MP or CRAY Y-MP processor. This speed advantage is coupled with an order of magnitude advantage in machine acquisition cost. SRIM is a somewhat large (30,000 lines of Fortran 77) program designed for uniprocessors; its restructuring for a hypercube provides new lessons in the task of alteri ...
- 17 Approximate query processing using wavelets** 77%
 Kaushik Chakrabarti , Minos Garofalakis , Rajeev Rastogi , Kyuseok Shim
The VLDB Journal — The International Journal on Very Large Data Bases

September 2001
Volume 10 Issue 2-3

Approximate query processing has emerged as a cost-effective approach for dealing with the huge data volumes and stringent response-time requirements of today's decision support systems (DSS). Most work in this area, however, has so far been limited in its query processing scope, typically focusing on specific forms of aggregate queries. Furthermore, conventional approaches based on sampling or histograms appear to be inherently limited when it comes to approximating the results of complex queri ...

18 Poster session: A four-bit full adder implemented on fast SiGe FPGAs 77%

 with novel power control scheme

K. Zhou , M. Chu , C. You , J.-R. Guo , Channakeshav , J. Mayega , B. S. Goda , R. P. Kraft , J. F. McDonald

Proceedings of the 2003 ACM/SIGDA eleventh international symposium on Field programmable gate arrays February 2003

The low operating speed of current CMOS Field Programmable Gate Arrays (FPGAs), i.e., 10-220 MHz, has prevented their use in high-speed digital applications. With the advent of IBM Silicon Germanium (SiGe) 7HP technology, designers have been able to design FPGAs operating in the gigahertz range. This paper is going to elaborate on the implementation of a 4-bit ripple-carry full adder (FA) on the new SiGe FPGA with new architectures and a novel power management strategy. The 1-bit FA can be reali ...

19 Poster session: A high resolution diagnosis technique for open and short 77%

 defects in FPGA interconnects

Mehdi Baradaran Tahoori

Proceedings of the 2003 ACM/SIGDA eleventh international symposium on Field programmable gate arrays February 2003

A two-step diagnosis flow, coarse-grain and fine-grain, is presented in order to identify a faulty element in the FPGA interconnects. The fault models used for interconnect are open, resistive-open, and bridging fault. The coarse-grain phase identifies the faulty net, the routing between two consecutive sequential elements in the FPGA. This phase is performed by just post-processing tester results for the test configurations used for interconnect testing. During the fine-grain step, the faulty n ...

20 Poster session: Application-dependent testing of FPGAs for bridging 77%

 faults

Mehdi Baradaran Tahoori

Proceedings of the 2003 ACM/SIGDA eleventh international symposium on Field programmable gate arrays February 2003

A new technique is presented for testing for bridging faults in the interconnects of an arbitrary design implemented in an FPGA. The configuration of the routing resources used in the original design remains unchanged in the test configurations. Only the logic blocks used in the design are reprogrammed in order to implement single-term functions, logic functions with only one minterm or one maxterm. As shown by formal proofs, all activated faults are detected when single-term functions and appro ...

Results 1 - 20 of 67 short listing


Prev Page 1 2 3 4 
Next Page



> home > about > feedback > login

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: [optimal <and> range <and> presentation <and> hierarchy <and> truncate <and> divide <and> magnitude <and> limit]
Found 30 of 151,219 searched.

Search within Results



> Advanced Search

> Search Help/Tips

Sort by: Title Publication Publication Date Score Binder

Results 1 - 20 of 30 short listing

Prev Page 1 2 Next Page

- | | |
|--|-----|
| 1 Curriculum 68: Recommendations for academic programs in computer science: a report of the ACM curriculum committee on computer science
William F. Atchison , Samuel D. Conte , John W. Hamblen , Thomas E. Hull , Thomas A. Keenan , William B. Kehl , Edward J. McCluskey , Silvio O. Navarro , Werner C. Rheinboldt , Earl J. Schweppe , William Viavant , David M. Young
Communications of the ACM March 1968
Volume 11 Issue 3 | 80% |
| 2 The Quadtree and Related Hierarchical Data Structures
Hanan Samet
ACM Computing Surveys (CSUR) June 1984
Volume 16 Issue 2 | 80% |
| 3 Practical byzantine fault tolerance and proactive recovery
Miguel Castro , Barbara Liskov
ACM Transactions on Computer Systems (TOCS) November 2002
Volume 20 Issue 4
Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ... | 77% |
| 4 Gross motion planning—a survey
Yong K. Hwang , Narendra Ahuja | 77% |

**ACM Computing Surveys (CSUR) September 1992**

Volume 24 Issue 3

Motion planning is one of the most important areas of robotics research. The complexity of the motion-planning problem has hindered the development of practical algorithms. This paper surveys the work on gross-motion planning, including motion planners for point robots, rigid robots, and manipulators in stationary, time-varying, constrained, and movable-object environments. The general issues in motion planning are explained. Recent approaches and their performances are briefly described, a ...

5 Information systems outsourcing: a survey and analysis of the literature 77%

Jens Dibbern , Tim Goles , Rudy Hirschheim , Bandula Jayatilaka

ACM SIGMIS Database November 2004

Volume 35 Issue 4

In the last fifteen years, academic research on information systems (IS) outsourcing has evolved rapidly. Indeed the field of outsourcing research has grown so fast that there has been scant opportunity for the research community to take a collective breath, and complete a global assessment of research activities to date. This paper seeks to address this need by exploring and synthesizing the academic literature on IS outsourcing. It offers a roadmap of the IS outsourcing literature, highligh ...

6 Object-based and image-based object representations 77%

Hanan Samet

ACM Computing Surveys (CSUR) June 2004

Volume 36 Issue 2

An overview is presented of object-based and image-based representations of objects by their interiors. The representations are distinguished by the manner in which they can be used to answer two fundamental queries in database applications: (1) Feature query: given an object, determine its constituent cells (i.e., their locations in space). (2) Location query: given a cell (i.e., a location in space), determine the identity of the object (or objects) of which it is a member as well as the re ...

7 Topical web crawlers: Evaluating adaptive algorithms 77%

Filippo Menczer , Gautam Pant , Padmini Srinivasan

ACM Transactions on Internet Technology (TOIT) November 2004

Volume 4 Issue 4

Topical crawlers are increasingly seen as a way to address the scalability limitations of universal search engines, by distributing the crawling process across users, queries, or even client computers. The context available to such crawlers can guide the navigation of links with the goal of efficiently locating highly relevant target pages. We developed a framework to fairly evaluate topical crawling algorithms under a number of performance metrics. Such a framework is employed here to evalua ...

8 Statistical Timing Analysis for Intra-Die Process Variations with Spatial 77%**Correlations**

Aseem Agarwal , David Blaauw , Vladimir Zolotov

Proceedings of the 2003 IEEE/ACM international conference on Computer-aided design November 2003

Process variations have become a critical issue in performance verification of high-performance designs. We present a new, statistical timing analysis method that accounts for inter- and intra-dieprocess variations and their spatial correlations. Since

statistical timing analysis has an exponential run time complexity, we propose a method whereby a statistical bound on the probability distribution function of the exact circuit delay is computed with linear run time. First, we develop a model for repre ...

9 XML query processing I: Dynamic sample selection for approximate query processing 77%

 Brian Babcock , Surajit Chaudhuri , Gautam Das

Proceedings of the 2003 ACM SIGMOD international conference on Management of data June 2003

In decision support applications, the ability to provide fast approximate answers to aggregation queries is desirable. One commonly-used technique for approximate query answering is sampling. For many aggregation queries, appropriately constructed biased (non-uniform) samples can provide more accurate approximations than a uniform sample. The optimal type of bias, however, varies from query to query. In this paper, we describe an approximate query processing technique that dynamically constructs ...

10 Lexicon design using perfect hash functions 77%

 Nick Cercone , Max Krause , John Boates

ACM SIGSOC Bulletin , Proceedings of the joint conference on Easier and more productive use of computer systems. (Part - II): Human interface and the user interface - Volume 1981 May 1981

Volume 13 Issue 2-3

The research reported in this paper derives from the recent algorithm of Cichelli (1980) for computing machine-independent, minimal perfect hash functions of the form: hash value ≡ hash key length + associated value of the key's first letter + associated value of the key's last letter hash value ≡ hash key length + associated value of the key's first letter + associated value of the key's last letter A minimal perfect hash function is one which provides sin ...

11 The early history of COBOL 77%

 Jean E. Sammet

ACM SIGPLAN Notices , The first ACM SIGPLAN conference on History of programming languages January 1978

Volume 13 Issue 8

This paper discusses the early history of COBOL, starting with the May 1959 meeting in the Pentagon which established the Short Range Committee which defined the initial version of COBOL, and continuing through the creation of COBOL 61. The paper gives a detailed description of the committee activities leading to the publication of the first official version, namely COBOL 60. The major inputs to COBOL are discussed, and there is also a description of how and why some of the technical decisions ...

12 Scalable feature selection, classification and signature generation for organizing large text databases into hierarchical topic taxonomies 77%

 Soumen Chakrabarti , Byron Dom , Rakesh Agrawal , Prabhakar Raghavan

The VLDB Journal — The International Journal on Very Large Data Bases August 1998

Volume 7 Issue 3

We explore how to organize large text databases hierarchically by topic to aid better searching, browsing and filtering. Many corpora, such as internet directories, digital libraries, and patent databases are manually organized into topic hierarchies, also called *taxonomies*. Similar to indices for relational data, taxonomies make search and access more efficient. However, the exponential growth in the volume of on-line

textual information makes it nearly impossible to maintain such taxono ...

13 A radar simulation program for a 1024-processor hypercube 77%

 J. L. Gustafson , R. E. Benner , M. P. Sears , T. D. Sullivan

Proceedings of the 1989 ACM/IEEE conference on Supercomputing August 1989

We have developed a fast parallel version of an existing synthetic aperture radar (SAR) simulation program, SRIM. On a 1024-processor NCUBE hypercube it runs an order of magnitude faster than on a CRAY X-MP or CRAY Y-MP processor. This speed advantage is coupled with an order of magnitude advantage in machine acquisition cost. SRIM is a somewhat large (30,000 lines of Fortran 77) program designed for uniprocessors; its restructuring for a hypercube provides new lessons in the task of alteri ...

14 Approximate query processing using wavelets 77%

 Kaushik Chakrabarti , Minos Garofalakis , Rajeev Rastogi , Kyuseok Shim

The VLDB Journal — The International Journal on Very Large Data Bases

September 2001

Volume 10 Issue 2-3

Approximate query processing has emerged as a cost-effective approach for dealing with the huge data volumes and stringent response-time requirements of today's decision support systems (DSS). Most work in this area, however, has so far been limited in its query processing scope, typically focusing on specific forms of aggregate queries. Furthermore, conventional approaches based on sampling or histograms appear to be inherently limited when it comes to approximating the results of complex queri ...

15 Classification of textures using higher-order fractal dimensions 77%

 A. Ait-Kheddache

Proceedings of the first international conference on Industrial and engineering applications of artificial intelligence and expert systems - Volume 2 June 1988

16 Timekeeping in the memory system: predicting and optimizing memory behavior 77%



Zhigang Hu , Stefanos Kaxiras , Margaret Martonosi

ACM SIGARCH Computer Architecture News May 2002

Volume 30 Issue 2

Techniques for analyzing and improving memory referencing behavior continue to be important for achieving good overall program performance due to the ever-increasing performance gap between processors and main memory. This paper offers a fresh perspective on the problem of predicting and optimizing memory behavior. Namely, we show quantitatively the extent to which detailed timing characteristics of past memory reference events are strongly predictive of future program reference behavior. We pro ...

17 Tuning Strassen's matrix multiplication for memory efficiency 77%

 Mithuna Thottethodi , Siddhartha Chatterjee , Alvin R. Lebeck

Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)

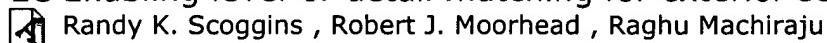
November 1998

Strassen's algorithm for matrix multiplication gains its lower arithmetic complexity at the expense of reduced locality of reference, which makes it challenging to implement the algorithm efficiently on a modern machine with a hierarchical memory system. We report on an implementation of this algorithm that uses several unconventional

techniques to make the algorithm memory-friendly. First, the algorithm internally uses a non-standard array layout known as Morton order that is based on a quad-tr ...

18 Enabling level-of-detail matching for exterior scene synthesis

77%



Randy K. Scoggins , Robert J. Moorhead , Raghu Machiraju

Proceedings of the conference on Visualization '00 October 2000**19 Overview of continuous optimization advances and applications to circuit tuning**

77%



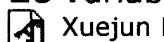
Andrew R. Conn , Chandu Visweswarah

Proceedings of the 2001 international symposium on Physical design April 2001

This paper surveys the state-of-the-art in continuous nonlinear optimization and makes the case that due to tremendous recent progress, larger and more complex problems can be solved than previously thought possible. The two basic paradigms, trust-region and line-search methods, are briefly described. In addition, various nonlinear optimization techniques are reviewed. The application of these nonlinear optimization methods to circuit sizing is presented by describing a pair of circuit si ...

20 Variable-precision rendering

77%



Xuejun Hao , Amitabh Varshney

Proceedings of the 2001 symposium on Interactive 3D graphics March 2001

Results 1 - 20 of 30 short listing **Prev
Page** **Next
Page**

1 2

The ACM Portal is published by the Association for Computing Machinery. Copyright ?2005 ACM, Inc.